



COPY OF PAPERS
ORIGINALLY FILED

RECEIVED

MAY 23 2002

Technology Center 2100

APPARATUS, METHOD AND COMPUTER SYSTEM FOR
REDUCING POWER CONSUMPTION OF A PROCESSOR OR
PROCESSORS UPON OCCURRENCE OF A FAILURE
CONDITION AFFECTING THE PROCESSOR OR PROCESSORS

BACKGROUND OF THE INVENTION

The present invention relates to a method and an apparatus for reducing power consumption of one or more processors in response to a failure condition affecting the processor or processors such as a thermal failure condition, and to a computer system incorporating such a method and/or apparatus.

In recent years, the processing power of processors has been increasing at a high rate. This increase in processing power has caused processors to heat up faster and to a higher temperature than previous processors. Therefore as the processing power has increased, a need has

RECEIVED

JAN 25 2002

OFFICE OF PETITIONS
DEPUTY A/C PATENTS

arisen for cooling the processor so that an overtemperature condition does not occur, since damage to the processor can occur if the temperature remains at too high of a level.

As the processing power of processors has increased and the need

- 5 for maintaining the processor at a relatively low temperature has become important, new ways of maintaining a relatively low temperature of the processor have been implemented. For example, heat sinks have been attached directly on or near the processor to help dissipate some of the heat from the processor. Additionally, cooling fans have been used to blow air in
- 10 the general vicinity of the processor or at the processor to help keep the processor from overheating. However, even when taking these measures, overtemperature problems can occur. Additionally, with the increase of processing power, the requirement for large heat sinks, blowers, cooling fans or other cooling mechanisms can cause size and expense problems,
- 15 and these mechanisms are still sometimes unable to properly cool some processors even when these cooling mechanisms are operating at full efficiency.

In case of a failure, the temperature of the processor can rise to an

overtemperature condition even if a cooling fan, a heat sink, or another

- 20 cooling mechanism normally is able to maintain the desired temperature of the processor. For example, the cooling fan may fail for some reason (i.e., the speed of the cooling fan may be reduced or it may completely stop). In

this case, the temperature of the processor can rise to a level which creates damage to the processor. Additionally, as processing power increases in current generation and future generation processors, additional measures for maintaining the temperature of the processor may become necessary.

- 5 Therefore, a need has arisen for additional ways of maintaining the temperature of the processor below a predetermined level. These additional measures may be in addition to or in place of current implementations using cooling devices such as heat sinks and cooling fans.

In addition to using cooling devices such as heat sinks and cooling

- 10 fans, other methods for ensuring that the temperature of a processor or processors does not become too high have previously been contemplated.

For example, a failure signal corresponding to a reduced performance of a cooling mechanism such as a cooling fan can be produced when the cooling mechanism either fails or has some sort of other reduction in performance

- 15 thereof. This signal is then used to completely shut down the processor, or provide a warning signal to the user of a personal computer or to a network manager, for example. However, if the signal is sent to the personal computer user or network manager (or other user) without turning off the processor, continued use of the processor could result in damage to the

- 20 processor or other components of the system. Similarly, if the processor is shut off, a resulting reduction in performance of the processor or processors occurs during the time which the processor is shut off. If the processor is

shut down, it is not operational until the failure is resolved. If such a processor is included in a uni-processor system, a system crash will occur and the entire system is shut down.

Additionally, other problems can occur with respect to signals sent to

- 5 the processor relating to vital functions of the system which are not received by the processor during the time which the processor is turned off. For example, the time-out of vital functions may occur if the processor is shut off for too long of a time period. These signals relating to vital functions of the system are sent to the processor for only a specific length of time before a
- 10 time-out of the signal occurs. If this time-out occurs, the processor does not receive the signals or perform any functions in response to these signals relating to vital functions of the system. For example, if a LAN (Local Area Network) network card is inserted and expecting a response from the processor, the network cards might drop clients if the processor does not
- 15 respond to certain signals prior to a time-out of those signals (i.e., within a predetermined time period). Therefore, a reduction of the power of a processor while still performing some processing functions would be beneficial so that the processor is not damaged and so that no vital functions of the processor or system are inadvertently not performed.

SUMMARY OF THE INVENTION

In accordance with the present invention, in order to reduce the power consumption of a processor, a failure signal is produced which indicates a failure condition affecting the processor. In response to the 5 failure signal, the power consumption of the processor is periodically reduced. This failure condition affecting the processor may be a thermal failure condition (or over-temperature condition).

In an illustrated embodiment, the processor includes a reduced power input. Power consumption of a processor is reduced in response to a first 10 signal level at the input and is not reduced in response to a second signal level at the input. A power reduction circuit provides the power reduction signal to the input of the processor in response to a failure condition affecting the processor. The signal provided to the input of the processor in response to the failure condition is a periodic signal alternately supplying the 15 first signal level which causes a reduction in the power consumption of the processor and the second signal level which does not cause a reduction in the power consumption of the processor.

The present invention allows a reduction in power consumption of a processor or processors in an economical manner when a failure condition 20 occurs. In described embodiments of the present invention, the failure condition is a thermal failure condition which occurs when the cooling mechanism fails or when a temperature of the processor or processors

increases to a high level. The present invention performs this reduction in power consumption without shutting down the processor entirely. The processor is allowed to continue to function at a reduced performance level without missing the receipt of signals provided to the processor relating to 5 functions of the system, which may include vital functions of the system.

In illustrated embodiments of the present invention, upon detection of a failure condition affecting the processor (such as a failure or a reduced performance of a cooling mechanism or a high temperature condition at or near the processor), the power consumption of the processor is periodically 10 reduced. This periodic reduction in power consumption may be implemented by periodically stopping and starting an internal clock of the processor or periodically reducing the power consumption of the processor in any other manner.

15 BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates an arrangement for reducing power consumption of a processor according to one embodiment of the present invention.

Figure 2 illustrates a periodic signal provided by the signal generator illustrated in Figure 1.

20 Figure 3 illustrates an arrangement for reducing power consumption of a processor according to an embodiment of the present invention.

Figure 4 illustrates an arrangement for reducing the power consumption of a processor according to an additional embodiment of the present invention.

Figure 5 illustrates an arrangement for reducing the power consumption of a processor according to a further embodiment of the present invention.

Figure 6 illustrates a multi-processor system in which the power consumption of one or more processors may be reduced according to an embodiment of the present invention.

Figure 7 illustrates an additional multi-processor system in which the power consumption of one or more processors may be reduced according to an embodiment of the present invention.

DETAILED DESCRIPTION

Figure 1 illustrates an arrangement for reducing the power consumption of a processor according to a first embodiment of the present invention. Illustrated in Figure 1 are a processor 10, a failure signal generator 12, a multiplexor 14 (MUX) and a signal generator 16. Processor 10 may be any processor or microprocessor including some processors of the microprocessor family developed by Intel Corporation commonly referred to as the x86 family of microprocessors (for example, the 80486, PentiumTM and Pentium ProTM microprocessors). Additionally, processor 10 may be a

later generation processor such as the Pentium Pro™ microprocessor or any other later generation processor.

In the illustrated embodiment, processor 10 includes a reduced power input 18. In response to an input signal at a predetermined level on this

5 input 18, the processor takes some action to reduce its power consumption.

For example, input 18 may be an input in response to which the internal clock of the processor 10 is stopped, thereby causing the processor to

consume less power. The Pentium™ and Pentium Pro™ processors include an input terminal for a STPCLK input signal which, when at a low level,

10 signifies a request to stop the internal clock of the processor and thereby cause the processor to consume less power. When the Pentium™ or

Pentium Pro™ processor recognizes the STPCLK input signal, the processor will stop execution on the next instruction boundary, unless superseded by a higher priority input, and generate a Stop Grant

15 Acknowledge cycle. When the STPCLK input signal is asserted, the

Pentium™ or Pentium Pro™ processor will still respond to external snoop requests. Although the STPCLK signal in the Pentium™ and Pentium Pro™

processors is a signal in which the active (or asserted) state occurs when the signal is at a low voltage level, in another processor it could be active at

20 a high level. Further, the processor need not have a STPCLK input. The present invention can be practiced using any processor input which stops

the internal clock of the processor, regardless of whether the active or

asserted state level is at the high or low voltage lev l. Some processors include inputs which are used to reduce the power consumption of the processor in some other way, or enable or disable the processor in some manner other than stopping the internal clock of the processor (e.g., a

5 processor enable or processor disable input signal). In such a case, these inputs may also be used in embodiments of the present invention.

Failure signal generator 12 has an output on line 20 which represents a presence of a failure condition affecting the processor 10. Examples of failure conditions affecting the processor include a failure of a cooling fan
10 cooling the processor or an overtemperature condition of the processor. However, according to an embodiment of the present invention, the signal on line 20 output by the failure signal generator 12 represents any failure condition affecting the processor 10.

Multiplexor 14 has two inputs A and B receiving signals on lines 20
15 and 22, a select line receiving a select signal on line 24, and one output coupled to the input 18 of the processor 10. As illustrated, the input on line 20 provided by the failure signal generator 12 is a signal representing a failure condition affecting the processor 10. The input on line 22 is a power reduction inactive signal corresponding to an inactive level of the power
20 reduction input signal 18. The input on line 24 is connected to an output of signal generator 16. In response to the failure signal 20, multiplexor 14

selects either the signal on line 22 or the signal on line 24 and outputs that signal on line 18.

In response to failure signal 20, the multiplexor 14 selects either the output 24 of signal generator 16 or a power reduction inactive signal level on

- 5 line 22 (in implementations using the Pentium™ and Pentium Pro™ processors or any other processors having a signal similar to the STPCLK signal, the power reduction inactive signal will be a high voltage level signal). In this manner, if failure signal 20 indicates that no failure condition affecting the performance of the processor 10 has occurred, the multiplexor
- 10 14 selects the power reduction inactive signal level 22 to be output to the power reduction input of the processor 10. In this case, the processor 10 operates without any reduction in power taking place. If the failure signal 20 indicates a failure condition of the processor 10, e.g., a reduction in power or failure of a cooling fan or sensed overtemperature of the processor 10, the
- 15 multiplexor 14 selects the output 24 of signal generator 16 to be provided to the power reduction input of processor 10. Multiplexor 14 may be any standard multiplexor or even a simple single pole double switch switching between inputs 22 and 24 based on the failure signal 20.

Signal generator 16 has two inputs on lines 26 and 28, respectively,

- 20 and provides an output on line 24. The input on line 26 is a signal representing a desired frequency and/or period and the input on line 28 is a signal representing a desired duty cycle. In response to inputs 26 and 28,

signal generator 16 provides a digital periodic signal on line 24 having on high value and one low value for each period.

Signal generator 16 is a standard well-known signal generator which generates a digital periodic signal on line 24 responsive to inputs 5 representing the desired frequency or period and/or duty cycle of the periodic signal to be provided to the power reduction input 18 of the processor 10. The duty cycle is defined as the ratio between the active level time period and the inactive level time period of the signal.

The frequency (or period) and duty cycle inputs to signal generator 10 10 may be predetermined values based on a variety of features of the system. For example, in determining the frequency (or period) input to the signal generator 16, a minimum processing interval and a maximum heat-up time must be considered. The minimum processing interval must be considered to ensure that all processing functions can be performed within the time 15 provided. The maximum heat-up time is a function of the die size, the maximum ambient temperature, and any thermal resistances to ambient temperature. The maximum heat-up time must be considered to reduce any thermal variation between the active and inactive signal levels. In determining the duty cycle, the minimum processing interval, a short enough 20 time to accommodate possible cooling failures such as fan failure, and an inactive time level which is short enough to ensure that no time-out conditions occur must all be considered. These considerations will vary

depending upon the system in which the processor 10 and other power reduction circuitry elements such as failure signal generator 12, multiplexor 14 and signal generator 16 are included.

Figure 2 illustrates an output signal from signal generator 16 which 5 may be used in implementing embodiments of the present invention such as the embodiment illustrated in Figure 1. The signal generator 16 generates a periodic signal including active signal levels 32 and inactive signal levels 34. The signal illustrated in Figure 2 can be applied as the STPCLK input to the Pentium™ and Pentium Pro™ processors, i.e., it is a low voltage level active 10 signal. Alternatively, the signal illustrated in Figure 2 would be inverted in an embodiment of the present invention in which the stop clock input to processor 10 were a high voltage level active signal.

The signal illustrated in Figure 2 can include a signal which may be output from signal generator 16 using the following system conditions: 15 Pentium™ processor and EISA refresh timeout of 100 μ sec. In such a system, a possible frequency of the signal illustrated in Figure 2 is 100 kHz. Additionally, a possible duty cycle of the signal illustrated in Figure 2 would be 25/75 (i.e., where the periodic signal has an active signal level 75% of the time and an inactive signal level 25% of the time). 20 While these values have been given as an example of the signal output from signal generator 16, any values providing a frequency and duty cycle meeting the following requirements may be used. Specifically, in

determining the frequency, the minimum processing interval of the processor and/or the entire system and the maximum heat-up time must be considered. In determining the duty cycle, the minimum processing interval, the ensuring of enough inactive time to accommodate possible cooling

5 failures, and an inactive time short enough to ensure that no time-out conditions occur should be considered. Additionally, any other input values may be used as inputs to the signal generator which may be used to describe a periodic signal (e.g., active and inactive level times, etc.)

Figure 3 illustrates an arrangement which may be used to reduce

10 power consumption of a processor according to another embodiment of the present invention. Figure 3 includes processor 10, multiplexor 14, signal generator 16 and cooling fan 36. Processor 10, multiplexor 14 and signal generator 16 function similarly to the corresponding elements of Figure 1. Therefore, a description of these elements is not included in the description

15 of Figure 2.

Cooling fan 36 is used to blow cool air in the direction of the processor 10. This cool air is used to maintain the processor 10 at a temperature low enough so that damage to the processor 10 does not occur due to an overtemperature condition. If reduced performance of cooling fan

20 36 occurs (for example, a reduction in the speed of the cooling fan or a failure of the cooling fan altogether), a fan failure signal 38 is provided. This fan failure signal 38 is provided to the select input of multiplexor 14. The fan

failure signal 38 may be directly provided from the cooling fan 36 or externally provided by a circuit detecting a failure or a reduced performance of the cooling fan 36. For example, an active level of the fan failure signal 38 may be provided if the speed of fan 12 falls below a predetermined level.

- 5 In the embodiment illustrated in Figure 3, the multiplexor selects the power reduction inactive signal on line 22 if the fan failure signal 38 indicates no fan failure and selects the output of signal generator 16 on line 24 if the fan failure signal 38 indicates a fan failure. The power consumption of processor 10 is therefore periodically reduced based on the fan failure
- 10 signal 38 the output of multiplexor 14 provided to the power reduction input 18 of the processor 10.

Figure 4 illustrates an arrangement which may be used to reduce the power consumption of a processor according to a further embodiment of the present invention. In Figure 4, processor 10 is mounted, for example, on a

- 15 printed circuit board (PCB) 40. A heat sink 42 is attached to processor 10 to provide an enhanced dissipation of heat from processor 10. A thermocouple 44 is embedded in heat sink 42 to measure a temperature near processor
- 20 10. Thermocouple 44 could be any temperature sensor device used to measure or sense the temperature at or near the processor 10 (such as a temperature sensor diode). Additionally, as an alternative to the thermocouple arrangement of Figure 4, embodiments of the present

invention may be practiced in which a device sensing any sort of failure at or near the processor 10 is used.

Thermocouple 44 provides an analog signal on line 46 to a connection 48 on the printed circuit board 40. This analog signal is 5 representative of the temperature at or near the processor 10. The analog temperature value is used to provide a failure signal similar to the failure signal 20 of Figure 1 to an arrangement similar to multiplexor 14 and signal generator 16, which provide an input to processor 10 used to reduce the power consumption of the processor in a manner similar to the arrangement 10 of Figure 1. An example of such an implementation is illustrated in Figure 5.

Figure 5 illustrates an arrangement according to an embodiment of the present invention which may be used in conjunction with the arrangement illustrated in Figure 4. Figure 5 includes a processor 10, a multiplexor 14, a signal generator 16, a temperature sensor 52, an analog- 15 to-digital (A/D) converter 54 and a look-up table 56.

Temperature sensor 52 senses a temperature at or near processor 10. Temperature sensor 52 may be a thermocouple located on or near processor 10 (e.g., thermocouple 44 of Figure 4) or any other temperature sensor. The sensed temperature output from temperature sensor 52 is 20 provided to analog-to-digital (A/D) converter 54. This signal may be provided, for example, via a connection such as connection 48 of Figure 4.

A/D converter 54 converts the sensed analog temperature to a digital signal representative of the temperature.

The digital signal output from A/D converter 54 and corresponding to the sensed temperature is provided to a look-up table 56. Look-up table 56

5 may be a Read Only Memory (ROM) or any other memory, for example.

Look-up table 56 stores values to be provided to the select input of

multiplexor 14 based on the sensed digital value. Each digital value has an

entry storing a corresponding value to be provided to the select input of

multiplexor 14. For example, any digital sensed values at or above a

10 predetermined temperature will reference an entry in the look-up table 56

corresponding to that digital value and storing a select signal for the

multiplexor 14 causing the multiplexor to select the output from the signal

generator 16. Similarly, in response to any value lower than the

predetermined temperature, the look-up table provides a select signal so

15 that the multiplexor 14 selects the power reduction inactive signal. In this

manner, the multiplexor 14 provides the periodic signal output by signal

generator 16 when the temperature at or near the processor 10 sensed by

temperature sensor 52 is at or above a predetermined temperature value

(e.g., 45°C or 85°C, etc.) and provides a power reduction inactive signal to

20 the power reduction input of processor 10 when the sensed temperature is

lower than that value. The predetermined temperature is preferably a

temperature well below the temperature at which the rated maximum wattage value of the processor will be reached.

As an alternative embodiment to the embodiment of Figure 5, a comparator can be used in place of the look-up table 56. The comparator

- 5 compares the temperature provided from A/D converter 54 with a predetermined temperature value and provides the result of the comparison to the select input of multiplexor 14.

Figure 6 illustrates an embodiment of the present invention in which a plurality of processors are arranged in a computer system. In Figure 6,

- 10 processors 60, 62 and 64 are included in a computer system such as a workstation or networking environment. Power reduction circuits 66, 68 and 70 are used to periodically reduce the power consumption of respective processors 60, 62 and 64 in response to a signal indicating a failure condition affecting the processor. Power reduction circuits 66, 68 and 70
- 15 may include the power reduction circuits illustrated in Figures 1, 3 and 4, for example. Specifically, power reduction circuits 66, 68 and 70 may each include a failure signal generator 12, a multiplexor 14 and a signal generator 16 as illustrated in Figure 1, or may include a cooling fan 36, a multiplexor 14 and a signal generator 16 as illustrated in Figure 3, or may include a
- 20 temperature sensor 52 (or thermo-couple 44), an analog-to-digital (A/D) converter 54, a look-up table 56, a multiplexor 14 and a signal generator 16 as illustrated in Figures 4 and 5. Each of the power reduction circuits 66, 68

and 70 may also be any other arrangement providing a signal used to periodically reduce the power consumption of a processor or processors. The power reduction circuits 66, 68 and 70 respectively provide a signal to the power reduction input of processors 60, 62 and 64 to periodically reduce

5 the power of the respective processor.

The input signal to the power reduction circuits corresponding to a reduction in performance of the processor could relate to a failure or reduction in performance of a cooling fan blowing air toward the respective processor or a temperature (or overtemperature condition) at or near the

10 respective processor.

Alternatively, a multi-processor embodiment of the present invention can be used in which one power reduction circuit provides the periodic signal to the power reduction inputs of all of the processors upon detection of a failure condition at or near any of the processors or of a failure condition

15 relating to an overall cooling mechanism such as a cooling fan which cools all of the processors, or any separate cooling fan corresponding to a particular processor. Such an embodiment is illustrated in Figure 7. Figure 7 includes a power reduction circuit 72 detecting a failure condition affecting one or more of the processors 60, 62 and 64 and providing a signal to the

20 power reduction inputs of one or more (or all) of the processors 60, 62 and 64 in response to the failure condition to periodically reduce power consumption of one or more (or all) of the processors 60, 62 and 64.

While the multi-processor embodiments of Figures 6 and 7 have been illustrated in a three processor system, it is noted that embodiments of the present invention may be implemented in systems including any multiple number of processors (i.e., two or more processors). That is, multi-
5 processor embodiments of the present invention are not limited to three processors.